

## CLAIMS

What is claimed is:

1. A method of optimizing storage of common data blocks within a networked storage system, said method comprising:
  - receiving a data block to be stored in a networked storage system;
  - analyzing contents of a received data block to determine a number of copies of the data block the networked storage system is entrusted to store and a number of copies of the data block existing within the networked storage system, and to identify a location of each copy of the data block within the networked storage system;
  - identifying performance and reliability requirements of the networked storage system;
  - determining an optimal number of copies of the received data block to store in the networked storage system, wherein the determining is made according to the number of copies of the data block the networked storage system is entrusted to store in combination with the identified performance and reliability requirements of the networked storage system; and
  - maintaining the optimal number of copies of the received data block within the networked storage system.
2. The method of claim 1, further comprising maintaining a record associated with each data block, wherein the record includes metadata specific to the data block for which the record is maintained.

3. The method of claim 1, wherein maintaining the optimal number of copies includes recognizing that a copy of the data block is no longer accessible, finding a duplicate copy of the data block, and generating another copy of the data block in order to achieve the optimal number of copies.
4. The method of claim 1, wherein the number of copies of the data block is based on data received in the networked storage system during a preceding predetermined period of time.
5. A method of optimizing storage of common data blocks within a networked storage system, said method comprising:
  - identifying a data block within a storage area network (SAN);
  - determining a number of copies of the data block existing within the SAN;
  - determining a number of copies of the data block the SAN is entrusted to store;
  - identifying performance and reliability requirements of the SAN;
  - determining an optimal number of copies of the data block to store in the SAN according to the number of copies of the data block the SAN is entrusted to store in combination with the identified performance and reliability requirements of the SAN; and
  - maintaining the optimal number of copies of the data block.

6. The method of claim 5, further comprising maintaining a record associated with each data block, wherein the record includes metadata specific to the data block for which the record is maintained.

7. The method of claim 5, wherein maintaining the optimal number of copies includes recognizing that a copy of the data block is no longer accessible, finding a duplicate copy of the data block, and generating another copy of the data block in order to achieve the optimal number of copies.

8. The method of claim 5, wherein the number of copies of the data block is based on data received in the SAN during a preceding predetermined period of time.

9. A storage system for optimizing storage of common data blocks within said storage system comprising:

storage blocks comprising data;

a file allocation table configured to the storage blocks and operable for matching a file identifier associated with a file to a location in the storage blocks allocated to the file;

a free space map operable for tracking unallocated storage blocks;

a redundancy management controller configured to the free space map and operable for maintaining an optimal number of copies of the data in the storage blocks; and

an occurrence count determination controller configured to the redundancy management controller and operable for determining a number of copies of the data in the storage blocks the storage system is entrusted to store.

10. The storage system of claim 9, further comprising a schedule control unit configured to the file allocation table and operable for scheduling operations of the occurrence count determination controller and the redundancy management controller at optimal times.

11. The storage system of claim 9, wherein the occurrence count determination controller is operable for determining an occurrence count of a block of the data by entering the data into a block contents map.

12. The storage system of claim 9, wherein the file allocation table comprises:  
a first data structure operable for mapping a file name to a file number; and  
a second data structure operable for mapping the file number to the storage blocks associated with the file.

13. The storage system of claim 9, wherein the free space map comprises a bitmap that marks a storage block as being allocated or unallocated.

14. The storage system of claim 9, wherein the number of copies of the storage blocks is based on data received in the storage system during a preceding predetermined period of time.

15. The storage system of claim 11, wherein the block contents map is operable for mapping the block of data to its possible locations in the storage system.
16. The storage system of claim 11, wherein the block contents map comprises a hash table.
17. The storage system of claim 11, wherein each entry of the data in the block contents map corresponds to a unique block of data in the system and comprises:
- an occurrence count for the block of data;
  - locations of storage blocks where the data is stored; and
  - for each of the storage blocks, file identifiers for files that have references to it and a number corresponding to the references.
18. The storage system of claim 17, wherein the redundancy management controller determines an amount of copies of the data to maintain based on a combination of the occurrence count for the block of data, performance settings of the storage system, reliability settings of the storage system, and attributes of the files.
19. The storage system of claim 11, further comprising a reverse block contents map, wherein the reverse block contents map is operable for mapping an address of the storage block to its corresponding block contents map entry.

20. The storage system of claim 10, wherein the schedule control unit schedules operations of the occurrence count determination controller and the redundancy management controller to be performed inline.

21. The storage system of claim 10, wherein the schedule control unit schedules operations of the occurrence count determination controller and the redundancy management controller to be performed offline at a time which reduces interference with processing of system operations comprising read, write, create, and delete operations received by the storage system.

22. The storage system of claim 9, wherein the storage blocks comprise data blocks each belonging to a single location in a single file.

23. The storage system of claim 9, wherein the storage blocks comprise data blocks each belonging to at least one location in at least one file.

24. The storage system of claim 9, wherein the storage blocks comprise unallocated storage blocks.

25. A networked storage system for optimizing storage of common data blocks within said networked storage system comprising:

means for identifying a data block within a storage area network (SAN);

means for determining a number of copies of the data block existing within the SAN;

means for determining a number of copies of the data block the SAN is entrusted to store;  
means for identifying performance and reliability requirements of the SAN;  
means for determining an optimal number of copies of the data block to store in the SAN  
according to the number of copies of the data block the SAN is entrusted to store in combination  
with the identified performance and reliability requirements of the SAN; and  
means for maintaining the optimal number of copies of the data block.

26. A program storage device readable by computer, tangibly embodying a program of instructions executable by said computer to perform a method of optimizing storage of common data blocks within a networked storage system, the method comprising:

receiving a data block to be stored in a networked storage system;  
analyzing contents of a received data block to determine a number of copies of the data block the networked storage system is entrusted to store and a number of copies of the data block existing within the networked storage system, and to identify a location of each copy of the data block within the networked storage system;  
identifying performance and reliability requirements of the networked storage system;  
determining an optimal number of copies of the received data block to store in the networked storage system, wherein said determining is made according to the number of copies of the data block the networked storage system is entrusted to store in combination with the identified performance and reliability requirements of the networked storage system; and  
maintaining the optimal number of copies of the received data block within the networked storage system.

27. The program storage device of claim 26, further comprising maintaining a record associated with each data block, wherein the record includes metadata specific to the data block for which the record is maintained.

28. The program storage device of claim 26, wherein maintaining the optimal number of copies includes recognizing that a copy of the data block is no longer accessible, finding a duplicate copy of the data block, and generating another copy of the data block in order to achieve the optimal number of copies.

29. The program storage device of claim 26, wherein the number of copies of the data block is based on data received in the networked storage system during a preceding predetermined period of time.

30. A program storage device readable by computer, tangibly embodying a program of instructions executable by said computer to perform a method of optimizing storage of common data blocks within a networked storage system, the method comprising:

identifying a data block within a storage area network (SAN);

determining a number of copies of the data block existing within the SAN;

determining a number of copies of the data block the SAN is entrusted to store;

identifying performance and reliability requirements of the SAN;

determining an optimal number of copies of the data block to store in the SAN according to the number of copies of the data block the SAN is entrusted to store in combination with the identified performance and reliability requirements of the SAN; and

maintaining the optimal number of copies of the data block.

31. The program storage device of claim 30, further comprising maintaining a record associated with each data block, wherein the record includes metadata specific to the data block for which the record is maintained.

32. The program storage device of claim 30, wherein maintaining the optimal number of copies includes recognizing that a copy of the data block is no longer accessible, finding a duplicate copy of the data block, and generating another copy of the data block in order to achieve the optimal number of copies.

33. The program storage device of claim 30, wherein the number of copies of the data block is based on data received in the SAN during a preceding predetermined period of time.